

# A Neural Network Approach to ECG Denoising

Rui Rodrigues (rapr@fct.unl.pt) and Paula Couto (mpcc@fct.unl.pt)

Departamento de Matemática da Faculdade de Ciencias e Tecnologia da UNL, 2829-516 Caparica, Portugal

December 19, 2012

## Abstract

We propose an ECG denoising method based on a feed forward neural network with three hidden layers. Particular useful for very noisy signals, this approach uses the available ECG channels to reconstruct a noisy channel. We tested the method, on all the records from Physionet MIT-BIH Arrhythmia Database, adding electrode motion artifact noise. This denoising method improved the performance of publicly available ECG analysis programs on noisy ECG signals. This is an offline method that can be used to remove noise from very corrupted Holter records.

## 1 Introduction

The ECG is often corrupted by different types of noise, namely, power line interference, electrode contact and motion artifacts, respiration, electrical activity of muscles in the vicinity of the electrodes and interference from other electronic devices. Analysis of noisy ECGs is difficult for humans and for computer programs. In this work we place ourselves in context of automatic and semi automatic ECG analysis: denoising should facilitate automatic ECG analysis.

General denoising signal processing methods have been applied to ECG. Low pass linear filters are used for high frequency noise removal, namely power line interference and muscle activity artifacts. High pass linear filters can be applied to cancel baseline wander. The use of neural networks to ECG denoising has been, to our knowledge, limited to the removal of these two types of noise. Other denoising tools are median filter, wavelet transform methods, empirical mode decomposition, morphological filters, non linear bayesian filtering and template matching. We will focus on noise introduced by electrode motion which causes more difficulties in ECG analysis[16]. Our method adapts to each particular ECG channel and learns how to reproduce it from a noisy version of the different channels available.

In the Physionet/Cinc Challenge 2010 it was shown that we can use some physiological signals to reconstruct another physiological signal, in particular an ECG [17, 20, 21]. This approach to reconstructing the noisy ECG channel is a simplified version, but equally effective, of the winning entry in that Challenge.

We show that the procedure is robust against noise in the input signals and can include, as an input, the channel we want to denoise.

This noise removal method is another example of the power of deep neural networks [8, 10, 5, 1], in this case, applied to ECG signals.

## 2 Method




Figure 1: Reconstructing the first channel from record 105 (MIT-BIH Arrhythmia Database), SNR=-6 db. In the lower section, noisy ECG, in the upper section, clean channel 1 and denoised channel 1.

If an ECG channel we want to use for ECG analysis is, at some time segment, contaminated with noise, we call it the **target** channel in our denoising process.

The method uses a feedforward neural network. A prerequisite for applying it is the target channel to be free from noise for some minutes, in order to train the neural network. The other channels used by the procedure may have noise in training time. If one channel has much more noise than others, it might be better not to use it for the reconstruction, even if it is the target channel.

The neural network will receive the samples as input, from one or more channels, corresponding to time segments with a fixed length. The output will be the samples from the target channel corresponding to the time segment used in the input. We used time segments with lengths of between one and three seconds, depending on the channels we use for reconstruction: one second if we use the target channel and another channel, two seconds, if we do not use the

target channel, and three seconds, if we only use the target channel.

To reconstruct one ECG channel we collected time segments,  $T_k$  from the ECG each one starting 16 samples after the preceding segment. After obtaining the output of the neural network, corresponding to each of the  $T_k$ , the value of the reconstruction on sample  $t_0$  will be the average value of the sample outputs corresponding to  $t_0$ , using all  $T_k$  that contain  $t_0$ .

The proposed method could be applied to a Holter record, reconstructing those time segments where an important level of noise is identified and using the remaining time of the Holter record for training.

## 2.1 Neural network architecture and training

We used a neural network with three hidden layers. The number of units on each of those layers was approximately 1000 in all experiments. To train the neural network, we constructed a sequence of time segments each one starting five samples after the beginning of the previous one. There is no need to use fiducial points to create input data to the neural network.

We applied Geoffrey Hinton’s method [9, 8, 6] to learn the neural network weights: following initialization using a stack of Restricted Boltzmann Machines, we applied backpropagation algorithm to fine tune the weights. For details on the training procedure for Restricted Boltzmann Machines, we refer to Hinton [7].

As usual, when using feedforward neural networks, we normalized the input data, to accelerate the learning process. First we applied a moving average filter, with the window size equal to the sampling rate. Then we subtracted the result from the signal, thus reducing the baseline wander. In the output signal, instead of the moving average filter we applied a median filter: it is more effective in the removal of baseline wander. Finally, we scaled the output signal to have unit variance and multiplied the input signals by the same scale factor.

We implemented our method using GNU Octave language and, to reduce training and reconstruction time, we ran most time consuming code on a Graphics Processing Unit. Our code is available at the first author’s web page.

## 2.2 Evaluating the method

Evaluating ECG denoising methods is not an obvious task. A common way of doing it is to add add noise to an existing signal and measure the Root Mean Square Error (RMSE) of the denoised signal relative to the original signal. This approach has some disadvantages. Firstly, when using a large data base of ECGs, is difficult to avoid noise in the original signal, and we do not want to punish the denoising method for not reconstructing the noise in the original signal. Secondly, RMSE does not always reflect the difficulties in analysing a noisy ECG. For instance, a constant baseline shift in the reconstructed signal is not very disturbing, but might correspond to a high RMSE value.

In this study we report RMSE in the reconstructed signal when we artificially add noise in the ECG, but we also evaluate our method using some publicly




Figure 2: Segment of first channel from record 202 (MIT-BIH Arrhythmia database): RMSE of reconstructed signal is largely due to the noise in the original signal: baseline shift and high frequency noise.

available programs that analyse the ECG: we compare the results of applying these programs with and without denoising the corrupted ECG. Although those programs already have a preprocessing stage to deal with noise, we show that, in the presence of noise, our denoising method improves their results.

### 2.2.1 Programs used to test this method

**gqrs** is a recent QRS detector, not yet published: the author is George Moody. This program is open source and available with WFDB software, from Physionet. There is an accompanying post-processor 'gqpost' intended to improve positive predictivity, at a cost of reduced sensitivity. We report the results of 'gqrs' with and without using gqpost. 'gqpost' uses a configuration file 'gqrs.conf'; we kept the default values of 'gqrs.conf'. The results of this program 'gqrs' depend on the value of a threshold parameter; as we did not find a systematic way of determining, for each record, the best value for the threshold, we used the parameter's default value. For this reason, we do not report the best possible results of this detector in the different records and therefore we should not use this study to compare the performance of the different qrs detectors.

**E.P. limited** is an open source program written by Patrick S. Hamilton [4]. It performs QRS detection and classifies beats as 'normal' or 'ventricular

ectopic' (VEB).

**ecgpuwave** It is an open source QRS detector and waveform limit locator, available as part of the PhysioToolkit [13, 12]. The authors are Pablo Laguna, Raimon Jan, Eudald Bogatell, and David Vigo Anglada.

All the programs listed above act on a single ECG channel, we did not find publicly available methods using more than one channel.

### 2.3 Statistics used to describe the results of QRS detectors and the beat classifier

For QRS detectors we used the following statistics:

$$\text{Sensitivity} = \frac{TP}{TP + FN}, \quad \text{Positive Predictivity} = \frac{TP}{TP + FP}$$

$$\text{Error rate} = \frac{FP + FN}{TP + FN}, [4]$$

where  $TP$  is the number of correctly detected beats,  $FP$  is the number of false detections and  $FN$  is the number of missed beats. For the beat classifier, we use, as in [11], the Sensitivity, Positive Predictivity,

$$\text{False Positive Rate} = \frac{FP}{TN + FP}$$

$$\text{and Classification Rate} = \frac{TN + TP}{TN + TP + FN + FP}$$

where  $TP$ ,  $TP$ ,  $FP$  and  $FN$  are defined as follows:

- **TP** is the number of beats correctly classified as VEB.
- **TN** is the number of non VEBs correctly classified.
- **FP** is the number of beats wrongly classified as VEB, excluding fusion and unclassifiable beats.
- **FN** is the number of true VEB not classified as such.

### 2.4 Adding noise to an existing ECG

In most experiments, to test the behavior of our denoising method, we start with a 'clean' ECG and add noise to it. For this we use the program **nst**, written by Geoge Moody [16].

The standard definition of signal to noise ratio (SNR), in decibels, is:

$$SNR = 10 \log_{10} \frac{S}{R}$$

where  $S$  and  $R$  are the power of signal and noise. We used a slightly different value for  $S$  and  $R$ , following the method used by the program 'nst'. Next we quote 'nst' man page [18]:

*“ A measurement based on mean squared amplitude, for example, will be proportional to the square of the heart rate. Such a measurement bears little relationship to a detector’s ability to locate QRS complexes, which is typically related to the size of the QRS complex. A less significant problem is that unweighted measurements of noise power are likely to overestimate the importance of very low frequency noise, which is both common and (usually) not troublesome for detectors. In view of these issues, nst defines  $S$  as a function of the QRS amplitude, and  $N$  as a frequency-weighted noise power measurement. “*

More details on the way ‘nst’ computes SNR can be found on the man page of ‘nst’.

### 3 Experiments

#### 3.1 MIT-BIH Arrhythmia Database




Figure 3: Reconstructing channel 1 from record 103 using only the same channel, SNR=0 db. At the bottom, noisy signal, in the middle denoised signal and at the top, clean signal.

We added noise to both channels in all the 48 records from Physionet MIT-BIH Arrhythmia Database [15, 3], and applied our method to reconstruct the first channel. As it is well known [16, 2], from the three types of noise, baseline

wander, muscle artifact and electrode motion artifact, it is the last one that creates most difficulties to ECG analysis programs. We contaminated both channels of each record with electrode motion artifact noise, using the corresponding noise record from the MITBIH Noise Stress Test Database [16].

In all but one record, both channels were used as input to reconstruct the first channel. In record 103, the noise in the second channel is already very high, therefore, we chose to use only the target channel in the reconstruction.

The clean record and corrupted noisy versions of the same record were used as input for training the neural network. We always used the clean target channel for the output.

The default behavior of the program 'nst' was followed to add noise to the records used in the tests: starting after the first five minutes from the beginning of each record, we added noise for two minutes, followed by two minutes without noise, repeating the process until the end of the record.

In order to train the neural network, we used all those segments of time where noise was not added in test time. In this way, the parts of noise used during training and testing do not overlap: we kept the neural network from learning the noise used in the test. The amount of noise, used for testing, corresponds to SNR values of 24 db, 18 db, 12 db, 6 db, 0 db and -6 db.

Table 1: Reconstruction Error: for each level of noise, we report the value of  $\text{RMSE}(\text{denoised signal})/\text{RMSE}(\text{noisy signal})$  over the 48 records of MIT-BIH Arrhythmia Database

SNR	24db	18db	12db	6db	0db	-6db
	1.014	0.507	0.257	0.133	0.071	0.041

In table 1 we report the fraction of RMSE, in the noisy signal, present in the reconstructed signal:  $\text{RMSE}(\text{denoised signal})/\text{RMSE}(\text{noisy signal})$ . As we can see in the table, there are no visible advantages, in terms of RMSE, in applying the denoising methods for low noise (SNR=24db), in fact, situations like the one in figure 2 introduce high values of RMSE because the method is not intended to learn to reproduce the noise of the original signal but just its main pattern. When the values of added noise increase, the errors in the reconstructed signal, due to noise in the original signal, lose their relative importance: for higher values of noise we notice an important reduction in the value of RMSE in the reconstructed signal. As supplementary material to this article, we present the detailed results for each record.

We applied the programs 'gqrs' and 'EPlimited' to the first channel, in noisy versions of each record and in the reconstructed signal, to verify whether, after applying our method, there were significant improvements in the performance of those programs. The results are reported in tables 2, 3, 4, 5. The first column indicates the SNR of the resulting ECG, the same value for both channels, after corrupting it with noise. The second column refers to the signal used when applying the program to the first channel: **denoised** means the reconstructed noisy first channel, using our method. The tables present the sensitivity, posi-

Table 2: Results of 'gqrs' applied to the first channel, corrupted with noise and denoised using our method (MIT-BIH Arrhythmia, 48 records).

SNR	channel 1	sensitivity	positive predictivity	error rate
24 db	noisy	0.9973	0.9963	0.0064
	<b>denoised</b>	0.9963	0.9992	0.0045
18 db	noisy	0.9973	0.9834	0.0195
	<b>denoised</b>	0.9964	0.9991	0.0045
12 db	noisy	0.9957	0.9054	0.1084
	<b>denoised</b>	0.9959	0.9991	0.0049
6 db	noisy	0.9881	0.7605	0.3231
	<b>denoised</b>	0.9941	0.9989	0.0071
0 db	noisy	0.9680	0.6471	0.5599
	<b>denoised</b>	0.9826	0.9922	0.0251
-6 db	noisy	0.9523	0.5806	0.7357
	<b>denoised</b>	0.9470	0.9466	0.1064

Table 3: Results of 'gqrs' with post-processor 'gqpost' applied to the first channel, corrupted with noise and denoised using our method (MIT-BIH Arrhythmia, 48 records).

SNR	channel 1	sensitivity	positive predictivity	error rate
24 db	noisy	0.9970	0.9970	0.0060
	<b>denoised</b>	0.9961	0.9994	0.0045
18 db	noisy	0.9969	0.9868	0.0165
	<b>denoised</b>	0.9960	0.9993	0.0047
12 db	noisy	0.9921	0.9318	0.0806
	<b>denoised</b>	0.9957	0.9992	0.0050
6 db	noisy	0.9659	0.8282	0.2345
	<b>denoised</b>	0.9942	0.9988	0.0069
0 db	noisy	0.9121	0.7037	0.4720
	<b>denoised</b>	0.9838	0.9947	0.0215
-6 db	noisy	0.8767	0.6301	0.6380
	<b>denoised</b>	0.9493	0.9568	0.0935

tive predictivity, number of detection errors and error rate, in the case of QRS detectors, and VEB sensitivity, positive predictivity, false positive rate and classification rate, for the 'EP Limited' beats classification. We used the following programs to report the results: 'bxb', from WFDB software [14], in the case of 'gqrs', and 'bxbep', in the the case of 'EPLimited'. The numbers are relative to all the 48 records from the MIT-BIH Arrhythmia Database, 91225 beats, from which 6101 are VEBs: we started the test after the first 5 minutes and stopped one second before the end. In the case of EP limited we had to start the test



one second later because we could not configure 'bxbep' to behave in another way.

Table 4: Results of EP Limited (QRS detection) applied to the first channel, corrupted with noise and denoised using our method(MIT-BIH Arrhythmia, 48 records).

SNR	channel 1	sensitivity	positive predictivity	error rate
24 db	noisy	0.9977	0.9981	0.0042
	<b>denoised</b>	0.9961	0.9996	0.0043
18 db	noisy	0.9977	0.9945	0.0079
	<b>denoised</b>	0.9957	0.9995	0.0048
12 db	noisy	0.9969	0.9342	0.0733
	<b>denoised</b>	0.9955	0.9995	0.0050
6 db	noisy	0.9857	0.7943	0.2696
	<b>denoised</b>	0.9944	0.9993	0.0063
0 db	noisy	0.9432	0.7110	0.4401
	<b>denoised</b>	0.9865	0.9951	0.0184
-6 db	noisy	0.8557	0.6568	0.5915
	<b>denoised</b>	0.9531	0.9493	0.0977

Table 5: Beat clasification as Normal or Ventricular Ectopic Beat (VEB): results of EP Limited applied to the first channel, corrupted with noise and denoised using our method(MIT-BIH Arrhythmia, 48 records).

SNR	signal	VEB sensitivity	VEB positive predictivity	VEB false positive rate	classification rate
24 db	noisy	0.9147	0.9589	0.0028	0.9916
	denoised	0.9142	0.9815	0.0013	0.9930
18 db	noisy	0.8873	0.9260	0.0051	0.9876
	denoised	0.9089	0.9759	0.0016	0.9923
12 db	noisy	0.8190	0.5935	0.0380	0.9530
	denoised	0.9032	0.9785	0.0014	0.9921
6 db	noisy	0.6977	0.2369	0.1291	0.8615
	denoised	0.8901	0.9778	0.0015	0.9912
0 db	noisy	0.6083	0.1308	0.2160	0.7751
	denoised	0.8567	0.9513	0.0032	0.9873
-6 db	noisy	0.5720	0.0939	0.2994	0.6940
	denoised	0.7663	0.7689	0.0166	0.9688

For QRS detectors, after applying our denoising procedure, there is always an improvement in positive predictivity in the tested programs, but, for high values of SNR there is a small reduction in the sensitivity: above 12 db for gqrs and above 6 db for EPLimited. Besides some ectopic beats not being well

reconstructed, the reduction in sensitivity is due to a smaller amplitude of the QRS complex in the reconstructed signal; this occurs in the first beat from figure 3. We could improve the sensitivity in the reconstructed signal, at the cost of a reduction in the positive predictivity, multiplying the reconstructed signal by a factor greater than 1.0, but we chose not to do it.

For beat classification there is always a clear improvement after using the proposed method.

As supplementary material to this article, we present the detailed results for each program and record.

### 3.2 Record mgh124

The MGH/MF Waveform Database [22, 3] is a collection of electronic recordings of hemodynamic and electrocardiographic waveforms. Typical records include three ECG leads. In the case of record mgh124, the first two ECG channels are sometimes strongly contaminated with noise, while the third ECG channel maintains a relatively good quality, therefore we have reliable QRS annotations. Using record mgh124, we tested our denoising method on a real ECG, without having to artificially add noise. In this case we reconstructed the second ECG channel, using only that same channel as input: we trained a neural network to produce a clean segment of the second channel given a corrupted version of the same segment. The clean parts of the channel 2 were used to obtain training data for the neural network.

tables 6 and 7 show the results of 'gqrs' and 'EP Limited' on the original second ECG channel and on the reconstructed version. The total number of beats during testing time is 8573, from which 458 are classified as VEBs.

Table 6: Results of 'gqrs' applied to second ECG channel, corrupted with noise and denoised using our method (record mgh124).

lead 2	detection errors	error rate	sensitivity	positive predictivity
noisy	1344	0.1568	85.06%	99.14%
<b>denoised</b>	890	0.1038	0.8971	0.9990

### 3.3 record sele0106 from QT database

Determination of peaks and limits of ECG waves is very important for ECG analysis: they are necessary for ECG measurements that are clinically relevant, namely, PQ interval, QRS duration, ST segment and QT interval.

Physionet QT data base was created to test QT detection algorithms [19]. Each record contains at least 30 beats with manual annotations identifying the beginning, peak and end of the P-wave, the beginning and end of the QRS-complex and the peak and end of the T-wave.




Figure 4: Reconstruction of noisy channel 2 from record mgh124: channel 3 is shown as a reference, it is not used for denoising. We notice that although the first VEB is recognized as such, the next two are not. This ECG segment starts at sample 853318.

Table 7: EP Limited: QRS detection and beat clasification as Normal or Ventricular Ectopic Beat (VEB), results on lead 2, corrupted with noise and denoised using our method(record mgh124).

signal	QRS detection errors	QRS error rate	QRS sensitivity	QRS positive predictivity
noisy	553	0.0645	0.9472	0.9878
<b>denoised</b>	463	0.0540	0.9470	0.9989

signal	VEB sensitivity	VEB positive predictivity	VEB false positive rate
noisy	23.14%	17.82%	6.305%
recsignal	36.68%	87.50%	0.311%

We used the program 'ecgpuwave' to show that, in some situations, we can improve automatic ECG delineation by using a clean channel to reconstruct a very noisy one.

Typically, the accuracy of ecgpuwave when detecting the limits or peak of some ECG characteristic wave is better in one of the channels. The best channel to locate one of those reference points changes with the different characteristic

points and also from record to record.

Table 8 shows the results of ecgpuwave, on the two channels, when it locates P wave peak, P ending, QRS beginning and QRS ending. We are using the first annotator as reference. We can conclude that the error is smaller when ecgpuwave is applied to the second channel.

Table 8: Accuracy of ecgpuwave, in ms, locating some characteristic points on the record sele0106(QT database): comparing results for the two leads.

reference point	channel	mean error	std error
P peak	1	9.33	4.54
	2	2.93	2.72
P off	1	15.73	10.06
	2	7.47	9.28
QRS on	1	18.80	6.62
	2	7.87	4.32
QRS off	1	13.33	4.17
	2	4.67	4.27

At this point we consider an easily imaginable situation, where the second channel is highly corrupted with noise, in such a way that it is better to use only the first channel for the reconstruction of the second channel. In this case we trained a neural network to produce a segment of channel 2 when it gets the corresponding segment of channel 1 as input.




Figure 5: reconstructing channel 2 from record sele0106(QT database) using only channel 1.

We followed this procedure and applied ecgpuwave to the reconstructed channel 2. The results are in table 9. One can see that we still get better results using reconstructed channel 2, from channel 1, than when applying ecgpuwave

to clean channel 1.

Table 9: comparing results, in ms, of ecgpuwave for the channel 1 and reconstructed channel 2

reference point	channel	mean error	std error
P peak	1	9.33	4.54
	reconstructed channel 2	2.80	2.76
P off	1	15.73	10.06
	reconstructed channel 2	7.20	6.54
QRS on	1	18.80	6.62
	reconstructed channel 2	7.33	4.63
QRS off	1	13.33	4.17
	reconstructed channel 2	6.40	4.21

## 4 Discussion of results and conclusions

Adding noise to existing records, we carried out extensive experiments on all the records from the MIT-BIH Arrhythmia Database. In the presence of high noise, SNR equal to 12db and lower, the programs we tested showed much better performance when we applied our denoising method to the ECGs. For low noise, SNR above 12db, after applying our method, QRS detectors show a slight reduction in sensitivity although there is an improvement in the positive predictivity. The experiments with records mgh124 and sele0106, without artificially adding noise in the test, confirm the advantages of using our method on a real ECG, a Holter record, for example. The experiment with record sele0106 also shows that the result of reconstructing a noisy channel can be exceptionally good when clean channels are available.

## References

- [1] Yoshua Bengio and Yann Lecun. Scaling learning algorithms towards ai. In L. Bottou, O. Chapelle, D. Decoste, and J. Weston, editors, *Large-Scale Kernel Machines*. MIT Press, 2007.
- [2] Gartheeban Ganeshapillai and John V Guttag. Real Time Reconstruction of Multi Parameter Physiological Signals. *EURASIP Journal on Advances in Signal Processing*, 2012(173), 2012.
- [3] A. L. Goldberger, L. A. N. Amaral, L. Glass, J. M. Hausdorff, P. Ch. Ivanov, R. G. Mark, J. E. Mietus, G. B. Moody, C.-K. Peng, and H. E. Stanley. PhysioBank, PhysioToolkit, and PhysioNet: Components of a new research resource for complex physiologic signals. *Circulation*, 101(23):e215–e220, 2000 (June 13). Circulation Electronic Pages: <http://circ.ahajournals.org/cgi/content/full/101/23/e215>.

- [4] P. S. Hamilton and W. J. Tompkins. Quantitative investigation of QRS detection rules using the MIT/BIH arrhythmia database. *IEEE Trans. Biomed. Eng.*, BME-33:1158–1165, 1987.
- [5] G. E. Hinton. To recognize shapes, first learn to generate images. In T. Drew P. Cisek and J. Kalaska (Eds.), editors, *Computational Neuroscience: Theoretical Insights into Brain Function*, chapter 3, pages 400–402. Elsevier, City, State of Publication, 1983.
- [6] G. E. Hinton. Learning Multiple Layers of Representations. *Trends in cognitive Sciences*, 11:428–434, 2007.
- [7] G. E. Hinton. A Practical Guide to Training Restricted Boltzmann Machines. Technical Report UTML TR 2010003, Dept. of Comp. Sci., University of Toronto, 2010.
- [8] G. E. Hinton. Reducing the dimensionality of data with neural networks. *Science*, 313(5785):504–507, July 2006.
- [9] Geoffrey E. Hinton. Learning to represent visual input. *Phil. Trans. R. Soc. B*, 365:177–184, 2010.
- [10] Geoffrey E. Hinton, Simon Osindero, and Yee-Whye Teh. A Fast learning Algorithm for Deep Belief Nets. *Neural Comput.*, 18(7):1527–1554, 2006.
- [11] Yu Hen Hu, Surekha Palreddy, and Willis J. Tompkins. A Patient-Adaptable ECG Beat Classifier Using a Mixture of Experts Approach. *IEEE Transactions On Biomedical Engineering*, 44(9):891–900, 1997.
- [12] R. Jané, A. Blasi, J. García, and P. Laguna. Evaluation of an automatic threshold based detector of waveform limits in Holter ECG with the QT database. *Computers in Cardiology*, 24:295–298, 1997.
- [13] P. Laguna, R. Jané, and P. Caminal. Automatic Detection of Wave Boundaries in Multilead ECG Signals: Validation with the CSE Database. *Computers and Biomedical Research*, 27(1):45–60, 1994.
- [14] G. B. Moody, C. L. Feldman, and J. J. Bailey. Standards and applicable databases for long-term ECG monitoring. *Journal of Electrocardiology*, 26:Suppl:151–155, 1993.
- [15] G. B. Moody and R. G. Mark. The impact of the MIT-BIH Arrhythmia Database. *IEEE Engineering in Medicine and Biology Magazine*, 20(3):45–50, 2001.
- [16] G. B. Moody, W. K. Muldrow, and R. G. Mark. A noise stress test for arrhythmia detectors. In *Computers in Cardiology*, pages 381–384, 1984.
- [17] George B. Moody. The PhysioNet/Computing in Cardiology Challenge 2010: Mind the Gap. In *Computing in Cardiology 2010*, volume 37, Belfast, 2010.

- [18] Moody, George B. nst man page. <http://www.physionet.org/physiotools/wag/nst-1.htm>.
- [19] Laguna P, Mark RG, Goldberger AL, and Moody GB. A Database for Evaluation of Algorithms for Measurement of QT and Other Waveform intervals in the ECG. In *Computers in Cardiology*, volume 24, pages 673–676, 1997.
- [20] R. Rodrigues. Filling in the Gap: a General Method using Neural Networks. In *Computers in Cardiology*, volume 37, pages 453–456, 2010.
- [21] A.M. Sullivan, H. Xia, and X. Zhao. Reconstruction of Missing Physiological Signals Using Artificial Neural Networks . In *Computers in Cardiology*, volume 37, pages 317–320, 2010.
- [22] J. P. Welch, P. J. Ford, R. S. Teplick, and R. M. Rubsamen. The Massachusetts General Hospital-Marquette Foundation Hemodynamic and Electrocardiographic Database – Comprehensive collection of critical care waveforms. *Journal of Clinical Monitoring*, 7(1):96–97, 1991.